Package-Level Interconnect Design for Optimum Electrical Performance

Lesley Polka, Assembly Technology Development, Intel Corporation
Shamala Chickamenahalli, Assembly Technology Development, Intel Corporation
Chee-Yee Chung, Assembly Technology Development, Intel Corporation
David G. Figueroa, Assembly Technology Development, Intel Corporation
Yuan-Liang Li, Assembly Technology Development, Intel Corporation
Kim Merley, Assembly Technology Development, Intel Corporation
Dustin Wood, Assembly Technology Development, Intel Corporation
Larry Zu, Optical Technology Group, Intel Capital

Index words: electrical interconnects, package design, electrical performance, power delivery, I/O, signal integrity, EMI

ABSTRACT

This paper presents an overview of the electrical analysis and electrical design processes and techniques used to develop package-level interconnect technologies to meet Intel's present and future performance challenges. The three areas of focus in electrical design of package interconnects: signal integrity, power delivery, and electromagnetic compatibility are discussed. Details of technology development and design choices made to meet the electrical performance requirements of some of Intel's most recent products are discussed. Some of the specific topics discussed include design and analysis flows, design techniques for impedance control and improved signal integrity, power decoupling solutions, voltage regulation module (VRM) analysis, and design for EMI control and improvement. Design and technology trade offs to meet the varied cost goals for different product lines resulting from Intel's recent market segmentation are also discussed using examples such as Intel's high-end server product line, Itanium™ processors, and Intel's value-PC line, Celeron™ processors. The importance of die-to-motherboard design integration early in the technology development cycle is also highlighted as critical in optimizing design for performance and cost.

INTRODUCTION

The last few years have seen significant improvements in microprocessor performance, and the expectation is that these performance trends will be maintained. Core speeds have surpassed the 1 GHz mark. Front-side bus speeds will soon go beyond 200 MHz; and many specialized buses operate at much higher transfer rates. At the same time, the environment in which this type of performance must be enabled has become increasingly challenging. Voltage levels have continued to decrease with each silicon technology generation to the point that < 1 V technologies are not too far out in the future. With tremendous increases in on-die transistor density, power levels have escalated exponentially to the point that they will soon surpass 100 W for many high-end microprocessors. This environment presents challenges for all aspects of electrical design of package-level interconnects. Lower voltage levels and faster bus speeds mean that noise and timing budgets are tighter. Higher power levels create challenges in power delivery through the package interconnect to the microprocessor core. Faster clock speeds and higher power levels create electromagnetic compatibility (EMC) and electromagnetic interference (EMI) challenges. Only a few years ago, the package-level interconnect was viewed as only a space transformer that bridged the gap between the fine silicon die features and the coarse features of the motherboard environment. The package provided mechanical support and protection for the fragile silicon die but was not viewed as an integral part of the electrical solution. This is no longer the case. In today's challenging microprocessor environment, a robust electrical design for the package-level interconnect can

enable superior system performance while a poor design can limit system performance.

Electrical design and analysis of a package and the next-level interconnects focuses on three main areas: I/O signal integrity, power delivery, and EMI control. Increased microprocessor core and bus speeds, decreased voltages, and increased power levels all present unique challenges to the package designer trying to ensure that electrical performance criteria are met. All of these areas require that the designer be very cognizant of the electrical characteristics and requirements of the rest of the system. The specific metrics and system requirements of interest for effective package-level interconnect design for I/O, power delivery, and EMI design are the subjects of this paper. The technology development and design processes for electrical package-level interconnects are presented. Each of the main areas of concern for package design: I/O, power delivery, and EMI is discussed in detail. A discussion of the design and technology tradeoffs necessary to meet Intel's market segmentation needs is also presented.

TECHNOLOGY DEVELOPMENT AND DESIGN FOR ELECTRICAL PERFORMANCE

Package Interconnect Overview

Package electrical design and analysis requires an awareness and understanding of the metrics used to gauge the electrical performance of the interconnect and of the manifestations of problems in interconnect designs. First, the physical aspects and metrics of an electronic package that are relevant to electrical performance need to be understood. Figure 1 shows a typical flip-chip package. To discuss the critical parts of the design, we separate them into the main areas of concern: I/O, power delivery, and EMI.

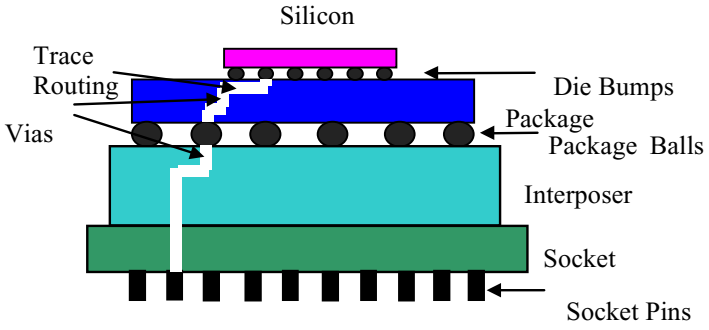

Figure 1: Flip-chip package

I/O Interconnect

The I/O interconnect refers to the connection between the driver and the receiver. This includes the chip connection (die bump or bondwire), vias, traces, and lands or pins in the package substrate. Structures analogous to those just described at the package level are included at the next levels of interconnect, i.e., the interposer, socket, cartridge, and motherboard. Similar interconnects on the receiver side of the system complete the I/O interconnect path. In addition, the return path must be identified and ultimately included as part of the I/O interconnect. The return path is the network through which current returns or leaves the die, depending upon the direction of the switching. It is, essentially, the nearest reference for the signal. If GTL logic is used, only a Vss reference needs to be nearby since GTL is referenced to Vss only. If CMOS logic is used, both a Vcc and a Vss reference need to be nearby; and one or the other is used as the reference depending upon whether the switching is from high-to-low or from low-to-high. Figure 2 roughly illustrates the optimal I/O interconnect structure in the package with the nearby return-path identified.

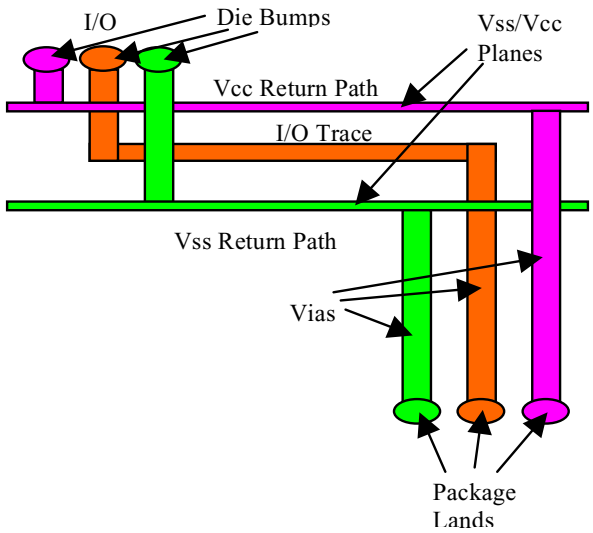

Figure 2: I/O package interconnect

Power Delivery Network

The package power delivery system is composed of the power and ground chip connection (die bumps or bondwires), power and ground vias, planes or strips for power and ground, and lands or pins in the package substrate dedicated to power and ground. Analogous structures exist in the interposer, socket, and motherboard. Discrete components that are integral to the power delivery network are discrete capacitors and the voltage regulator module (VRM). Figure 3 is a schematic of the power delivery network.

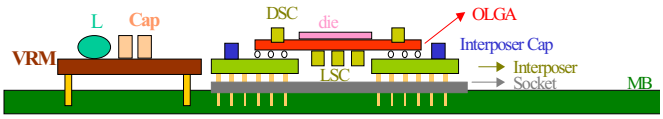

Figure 3: Power delivery network

Electromagnetic Interference (EMI)
For EMI design and analysis, the key features of interest include the location and routing of high-speed, repetitive signals, such as clock lines. The location of ground planes that can be used for shielding is also important. Grounding point location and quantity on any large conductive structures are also important to the design. For example, heat sinks, which are a critical aspect of the thermal solution, must be properly grounded or they will create EMI problems. Any materials that will be used as shielding materials are also part of the physical structure and need to be included in EMI design and characterization efforts.

Modeling and Design Overview

Good silicon models are also necessary for effective package design and technology development since the silicon model is used as the input in any package-level analysis, and the analysis results are very sensitive to small errors in the die models. In performing analyses to develop high-performance package solutions, it is necessary that the package developer/designer have access to die and system-level information. For signal integrity analysis, accurate buffer models are required. For power delivery analysis, accurate models of the on-die current switching signature and of the on-die decoupling capacitance are required. For EMI analysis, an understanding of the silicon switching and power hot spots is important.

As microprocessor performance reaches faster speeds, it also becomes more and more critical that package technology development and design are linked to silicon and system-level design activities early in the product definition stage. The reason is that margins are so tight that an integrated design is the only way to ensure optimized performance. Often trade offs or changes in the silicon or motherboard design or technology choice can greatly influence the effectiveness of the package-level solution or the ability to even find a solution that is cost effective and manufacturable. Although package-interconnect technology development and design focuses on the interconnect between the die and the motherboard, which consists of the package substrate, socket, and interposers or cartridges, knowledge of the entire system is becoming more and more critical to developing the most effective package-level interconnect technologies.

Key Metrics

The next critical area in package technology development and design for electrical performance is developing an understanding of the key metrics used to measure electrical performance. Again, it is best to discuss the metrics separately, i.e., to discuss each of the three areas of concern (signal integrity, power delivery, and EMI).

Signal Integrity
From a signal integrity, or I/O, perspective, the impact of the package interconnect on overall system performance is gauged primarily by the impact on timing criteria. The timing criteria are, in turn, impacted by signal integrity phenomena. Timing criteria and budgets are used to set signal integrity budgets and targets for each part of the I/O interconnect, including the package interconnect portion of the overall interconnect. The package designer uses the signal budgets and targets to guide the package design. The physical design of the package directly impacts the signal integrity characteristics of the package. This is a very high-level view of I/O interconnect design and analysis.

Power Delivery
The power delivery aspects of a package design are characterized by both DC and AC performance. The DC performance is measured in terms of the DC voltage drop through the package-level interconnect and in terms of the maximum current-carrying capability of the interconnect. The AC performance is measured in terms of di/dt, or switching noise, specifically in terms of noise on the core power supply measured at certain instances referred to as "1st droop," "2nd droop," and "3rd droop." Figure 4 illustrates the concept behind these terms. The 1st droop is generally mitigated by effective placement of high-frequency on-die and mid-frequency on-package decoupling. The 2nd droop is affected by package-level and low-frequency motherboard decoupling, and the 3rd droop is affected by motherboard decoupling and VRM placement. Since the voltage noise generated due to the di/dt switching is proportional to L di/dt, where L represents the power loop inductance, careful design of the power delivery network to mitigate this inductance is a critical aspect of power delivery design, especially at the package level. Much effort is expended in placing power and ground planes, power and ground vias, and in capacitor pad design to ensure low inductance power delivery loops.

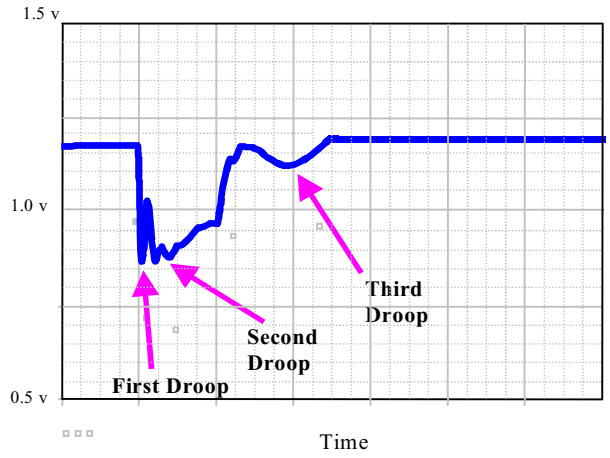

Figure 4: Power delivery di/dt noise and voltage-level droops

EMI

EMI is measured in terms of radiation levels given in decibels (dB). Since the Federal Communications Commission (FCC) regulates the amount of unwanted radiation that electronic devices can radiate across a frequency band related to the speed of the device, package-level EMI is usually gauged across a frequency band that extends well into the GHz region for today's microprocessors. Package-level EMI contributes to the overall EMI of a system; however, the allowable EMI from a package is not specified by FCC regulations. Instead, a budget is usually set for the amount of EMI suppression desired at the package level in order to effectively contain the EMI of the entire system. For example, grounding the thermal solution at the package level might amount to a 6 dB reduction in overall system-level EMI. Thus, package-level EMI solutions are usually characterized in terms of the amount of radiation reduction that they offer, but are not characterized in terms of absolute EMI levels. This is because the absolute level is a function of other elements besides the package-level solution, such as the die activity and the system-level EMI solution.

Analysis Overview

Package design is usually done iteratively. That is, budgets and targets are known up front, and a design is proposed. The design is then analyzed to see if it meets the budgets and targets. If not, it is tweaked; and more analysis is performed. Figure 5 illustrates this design process flow. Analysis during the design cycle consists of modeling and simulation using software tools and of validation using test vehicles. Modeling tools are used to generate circuit models of the package-level interconnect. Software packages that solve Maxwell's equations for the electromagnetic fields generated in the package interconnect are used to create appropriate models. These models are input into a SPICE circuit simulator, and simulations are run to investigate the performance of the package.

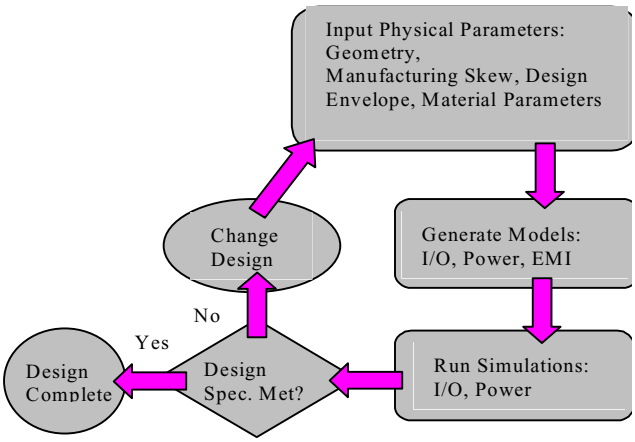

Figure 5: Electrical analysis and design flow

Signal Integrity

A general understanding of the types of analyses that are performed in designing a package to meet electrical performance requirements is critical to understanding the key package parameters and electrical characteristics of interest. Specifically, for I/O, or signal integrity, analysis, both single-bit and multiple-bit switching are investigated. "Single-bit switching" means that only one I/O buffer is switched, and the timing and signal integrity of the single I/O interconnect are observed. "Multiple-bit switching" means that many bits are switched simultaneously, and the timing and signal integrity impact is observed. A detailed understanding of system-level signal integrity analysis and design methodologies is not necessary to be able to extract the important design aspects of the package interconnect. A brief description of the metrics of interest is given in the next few paragraphs, however, and is useful in order to have a more complete understanding of the package design process.

Very basically, a signal takes a certain finite amount of time to travel from the driver to the receiver. This time is referred to as the "flight time." A signal is not considered to be "received," or "latched in," until it has been stable for a finite amount of time, which is called the "settling time." The concepts of flight time and settling time are illustrated in Figure 6.

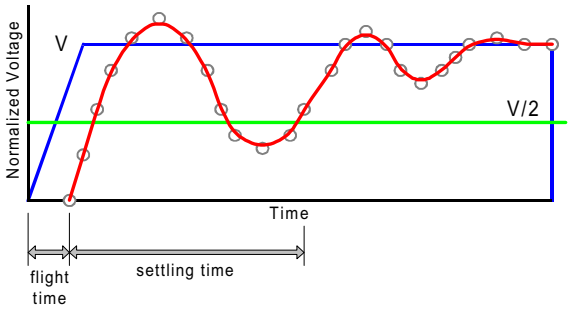

Figure 6: Flight time and settling time

Flight time is, very basically, the intrinsic delay through the interconnect and is affected by the physical interconnect length and the speed with which a signal can propagate through that distance. The velocity is indirectly proportional to dielectric constant (ε_r), which is an intrinsic property of the material surrounding the conductor through which the signal is traveling. Settling time is affected by many other phenomena, which will be described shortly, that occur along the I/O interconnect.

Figure 7 illustrates, very simplistically, the impact of the I/O interconnect on the signal. The upper graph illustrates the ideal driver output; the lower graph illustrates the receiver input, which differs quite a bit from the ideal due to the effects of the interconnect (i.e., between the driver and the receiver). The goal of I/O interconnect design is to ensure that the "signal integrity" of the driver output is maintained to an extent that enables the receiver to detect and to correctly interpret it.

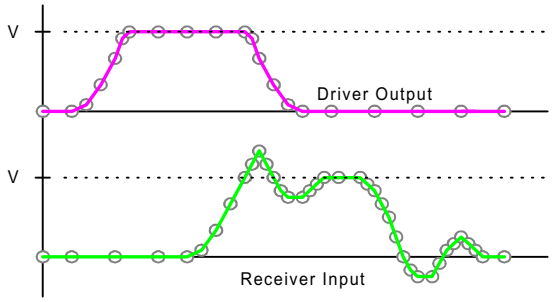

Figure 7: Signal integrity illustrated

A signal is considered "switched" or "received" after it has reached certain voltage thresholds that are defined according to the silicon technology and system architecture. The voltage drop between the driver and the receiver along the I/O interconnect due to the resistance of the interconnect directly impacts the signal's ability to reach the appropriate voltage level required for correct detection at the receiver.

Returning to the concept of settling time, a signal is not considered to be received until it is stable at the designated voltage levels for a fixed period of time called the settling time. Several signal integrity phenomena can occur at the receiver and degrade, or prolong, the settling time. Some of these concepts are illustrated in Figure 8. Ringback is when a signal continuously crosses the threshold voltage level before ultimately settling at that voltage level. Over/undershoot is when a signal transitions well beyond the threshold voltage before returning to that threshold. Nonmonotonicity is when a signal rings within the threshold voltage region. These signal integrity problems cannot be attributed to any one aspect of an interconnect design. Instead, they are the manifestation of any number of issues in the design. Some of the interconnect characteristics that directly impact the signal integrity of the design are the characteristic impedance of the signal line, discontinuities in the signal path, crosstalk between signal lines, and other phenomena that result when several signals switch simultaneously. These concepts are discussed in more detail in the next section of this paper.

Power Delivery

The goal of a good power delivery design is to maintain steady power and ground rails during core switching. The robustness of a power delivery design is measured by observing the power and ground levels during many cycles of on-die switching. Some deviation from the ideal power and ground levels is expected during switching. These deviations are referred to as voltage droops, and allowable budgets for each level of voltage droop are allotted. If the voltage droop goes beyond these budgets, the power delivery system is considered ineffective and design changes are necessary. The voltage droop levels are directly impacted by the loop inductance of the power delivery network as well as by the placement and quantity of capacitors and by the placement of the VRM. AC power delivery simulations are used to analyze and to determine the most effective design of the power delivery system.

EMI

EMI at the package level is directly impacted by the manner in which large metal areas are grounded and by the location of high-speed, repetitive signals, such as clock lines. Phenomena such as crosstalk, poor return-current paths and large levels of power/ground noise can impact EMI levels. These are also signal integrity and power delivery problems. In many instances, EMI problems can be traced to poor signal integrity and power delivery designs. Thorough EMI analysis, therefore, entails thorough signal integrity and power delivery analysis. In addition, EMI models of the package-level interconnect are used to directly look at the impact on radiated emissions of different EMI solutions, such as grounding schemes.

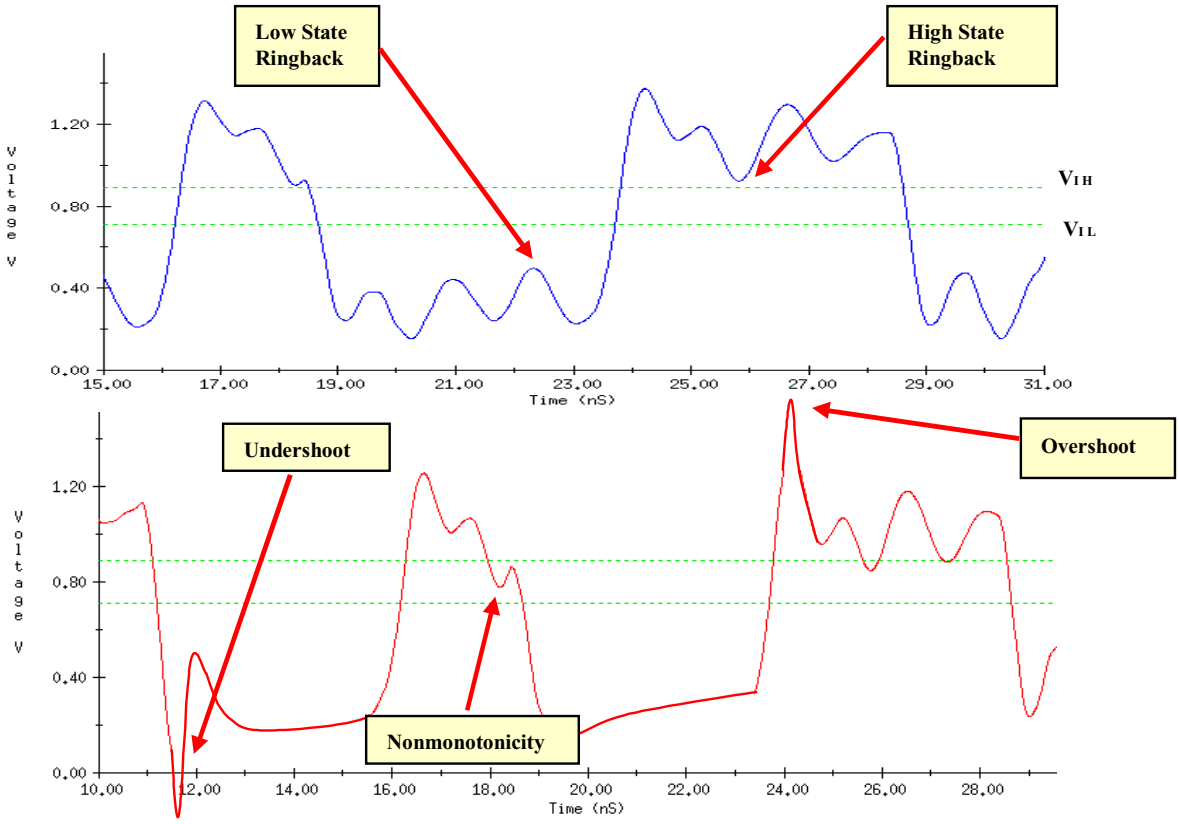

Figure 8: Signal integrity concepts illustrated

Design Flow

To summarize, system-level performance is measured in terms of timing, signal integrity, power-supply noise, and radiated emissions budgets. The impact of the package-level interconnect is directly attributable to electrical characteristics of the interconnect, namely, its characteristic impedance, crosstalk and simultaneous switching output (SSO) characteristics, resistance, loop inductance, and capacitance. These, in turn, are a direct function of the interconnect physical design and material selections. Characteristic impedance is determined by the package stack-up, the material dielectric constant, and the width and height of the trace. Resistance is determined by the resistivity of the trace metal, the cross-sectional area of the trace, and the trace length. Crosstalk is determined by the trace spacing, package stack-up, and dielectric constant of the material. SSO is a broad term used to refer to degradation in system performance due to effects seen when many signals switch simultaneously. The causes of SSO problems are many and varied. One of the package design choices that directly causes SSO problems is the inductance of the return path for each signal. This is directly related to the proximity of Vcc and Vss pins, planes and vias to the I/O path, and to discontinuities in the return path. Loop inductance of the power delivery network is impacted by the location and orientation of power and ground planes and vias and by the location and orientation of discrete capacitors.

The design flow and choices, thus, are very important in determining the interconnect's electrical characteristics and, ultimately, the performance of the package in the overall system. The design choices related to the package-level interconnect that must be made early in the design cycle are as follows:

- stack-up
- pinout
- pin pitch of the package, interposer, and socket
- trace routing topology (location of voids and degassing holes in the planes, amount of parallelism of traces, trace spacing, number and location of power/ground reference planes with respect to the routing layers, whether traces are routed as microstrips or striplines, etc.)
- min/max trace length and trace-length skews
- die bump pattern (location of power/ground bumps)
- capacitor location and connection
- VRM location and design
- grounding scheme for heat sinks and other metal structures

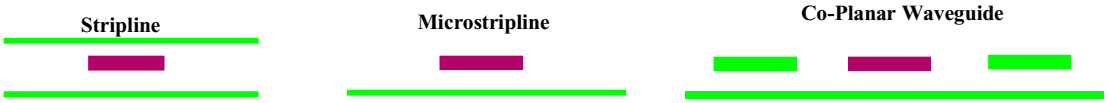

Figure 9: Package trace configurations

- material choices (dielectric constant and resistivity)

These choices directly impact the following electrical characteristics that are of primary concern in package interconnect design:

- characteristic impedance (Z_0)
- variation in characteristic impedance (ΔZ_0)
- resistance (R)
- mutual inductance (L_m) and capacitance (C_m) between traces (affects crosstalk)
- inductance of the return path ($L_{loop\text{-}I/O}$) (affects SSO)
- power delivery network loop inductance ($L_{loop\text{-}power}$)
- shielding effectiveness of EMI solution

These are discussed in more detail in the remainder of this paper.

SIGNAL INTEGRITY CHALLENGES AND SOLUTIONS

The challenge in designing a package with good signal integrity lies in properly selecting a technology and appropriately designing with that technology such that characteristic impedance, resistance, crosstalk, and SSO are controlled. These aspects of a package are related to fundamental choices such as the materials used. The ability to control the manufacturing process is also key to controlling these parameters within certain ranges. Finally, good design practices are necessary to ensure an optimized design that meets performance requirements.

Characteristic Impedance

Characteristic impedance is a term borrowed from transmission-line theory and is, at a very academic level, the ratio of the electric-to-magnetic field distribution in a transmission line, where a transmission line is a multi-conductor structure that propagates a signal [1]-[4]. Package and system designers are concerned about this characteristic of the I/O traces in the system because signal reflection occurs at any boundary where there is a change in characteristic impedance. The amount of reflection is directly proportional to the difference in impedances at the boundary according to

$$\Gamma = \frac{Z_1 - Z_0}{Z_1 + Z_0} \qquad (1)$$

where

Γ = the amount of reflection

Z_0 = the characteristic impedance of the trace

Z_1 = the characteristic impedance of the adjoining trace, connection, or media.

Reflection degrades signal integrity and is a primary cause of phenomena such as ringback and over/undershoot. The emphasis in I/O interconnect design is, therefore, on matching the impedance throughout the path of the I/O interconnect.

Although closed form solutions for characteristic impedance exist for certain configurations that are very similar to the designs used in traces in Intel's packages, these formulas are for aspect ratios of trace width to plane separation that are different from those of Intel's packages. The formulas, therefore, yield Z_0 values that are not very accurate in many cases. The preferred method of solving for Z_0 is to use a software tool that solves Maxwell's electromagnetic field equations. Z_0 can be measured in the lab using time-domain reflectometry (TDR). Frequency-dependent characteristics can be extracted using network analyzer measurements. Figure 9 illustrates the different designs used for traces in package substrates. If power/ground planes are above and below the trace, the structure is called a "stripline." If there is only one plane nearby, it is called a "microstripline." If the structure also contains shielding traces tied to ground on either side of the trace, it is called a "co-planar waveguide." The stripline configuration is the preferred configuration because the presence of two planes allows for better impedance control, provides higher bandwidth, and affords better isolation of signals from other parts of the package and system. Figure 10 shows a more detailed depiction of a stripline structure.

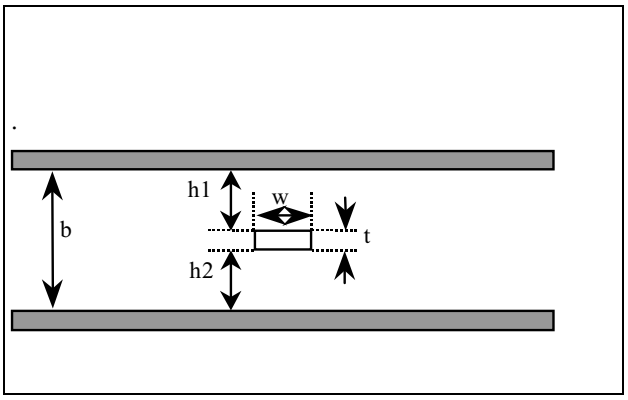

Figure 10: Stripline configuration

There are two items of concern in package design regarding Z_0. The first is being able to target a specific characteristic impedance, and the second is being able to control that impedance within a certain tolerance of the target value. The target value is usually set by matching to the system impedance, which is usually 50 ohms. Special bus protocols may require matching to an impedance other than 50 ohms. Sometimes matching to two or more different impedances within a package is required. The challenge in package design is, thus, to be able to target two or more impedances in a package and to be able to control that impedance within increasingly tighter tolerances. Variation in impedance is caused by manufacturing variations and differences in the topology of the routing in the package substrate. Since manufacturing variation will always be present, Intel has focused during package development on understanding the impact of the packaging technology and design on the target impedance and variation and on coming up with new designs to control these.

One example of a design innovation used to control impedance variation can be found in the packaging for the Pentium® II brand. In the original organic flip-chip packaging used for this product line, the power and ground planes above and below signal routing layers had degassing holes arranged periodically throughout the planes. These degassing holes were necessary for manufacturing reasons. The requirement was that there be a certain density of holes spread evenly throughout the planes, so routing above and beneath these holes was unavoidable. Originally, the degassing holes were aligned among all planes such that some traces were routed beneath solid metal and some beneath holes, as shown in Figure 11. This led to large impedance variation (~20%) due to routing variation. Investigations were made to develop an optimized degassing hole pattern such that the impedance variation would be minimized [5]. The optimized pattern is shown in Figure 12 and consists of a staggered, rotated pattern where any trace routed between these planes will "see" the same degassing hole pattern and, thus, have a very similar impedance to other traces routed on the same layer. Using this pattern, our analysis showed that the variation due to routing was reduced to <3%.

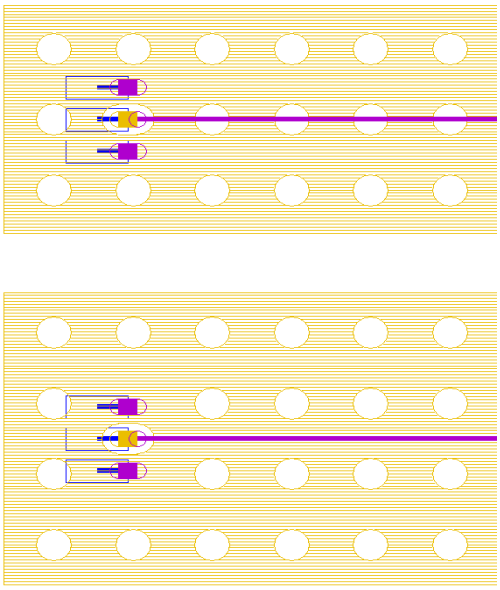

Figure 11: Old degassing-hole pattern

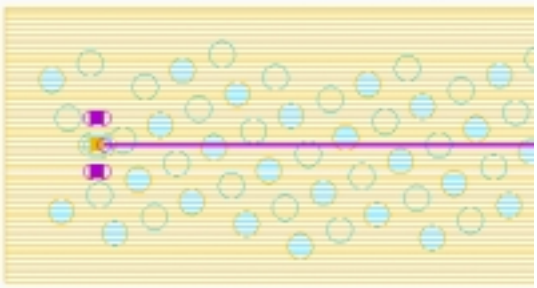

Figure 12: New degassing-hole pattern

Trace Resistance

After characteristic impedance, trace resistance is the next challenge faced in designing a package with good signal integrity. Resistance is a fairly easy concept to comprehend; yet, I/O package interconnect resistance has already been a concern on recent products and will continue to be an issue for future products. Resistance (R) of a structure is easily calculated by

$$R = \frac{\rho l}{A} \qquad (2)$$

where

ρ = the metal resistivity

l = the length of the structure parallel to the flow of current

A = the cross-sectional area through which current flows.

One of the challenges has been to obtain accurate values for ρ. The copper used for traces in Intel's packages is not pure copper; thus, the resistivity value is not easily obtainable from published literature. Much focus is given to obtaining accurate values for ρ during the development cycle by using test structures that are measured both internally and monitored continuously during process development at the suppliers. On top of the difficulty in obtaining accurate material parameter data for trace resistance, the inherent limitations of the materials used for package traces is a concern for future-generation package technologies. The reason is that trace widths are decreasing as die and package dimensions decrease, which increases the trace resistance. At the same time, voltage levels, and, thus, the allowable DC voltage drop through a package, are decreasing. These contradictory trends will prove to be a challenge that will need to be overcome in future package technology development.

Crosstalk

Impedance control and trace resistance are signal integrity characteristics of a package that are of concern for single-bit switching. The remaining signal integrity characteristics, crosstalk and SSO, are of concern for multiple-bit switching. Crosstalk is the term used to refer to unwanted signal coupling between lines. Crosstalk is a complex function of the characteristics of the driver and receiver, trace characteristics, and switching patterns; however, certain generalities regarding package design and its relationship to crosstalk can be made. Crosstalk in a package is dependent upon the stack-up, just as is characteristic impedance; however, crosstalk is primarily affected by the distance between traces, by the amount of parallelism between traces, and by the presence of reference plane discontinuities. The longer the parallel distance between two or more lines, the higher the crosstalk between them. Also, the closer the lines are to one another, the higher the crosstalk. The proximity of power/ground planes to the traces can help reduce crosstalk as long as these planes do not have discontinuities in the proximity of the traces. Reference plane discontinuities contribute to increased crosstalk levels for traces routed above or beneath these planes. For continuous reference planes, crosstalk is directly proportional to the distance between the planes above and below the trace.

There are no simple formulas for predicting crosstalk. Models of the traces, with the mutual inductance and capacitance calculated, must be generated and simulations run using buffer models and models representing the system loads to correctly determine the amount of crosstalk in a package design. There is a quick rule of thumb that can be used to approximate the amount of crosstalk expected in a package substrate or motherboard routed in a stripline configuration. Using the ratio of S/H, where S is the trace spacing and H is the distance from the trace to the power/ground plane above or below (assuming the traces are centered between the planes), the approximate crosstalk percentage is as follows [6]:

- if S/H > 1, crosstalk < 50%
- if S/H > 3, crosstalk < 10%
- if S/H > 4, crosstalk < 5%
- if S/H > 5, crosstalk < 1%

To meet today's product performance requirements, package-level crosstalk must be less than 10%, which means that S/H should be greater than 3. For today's flip-chip technology, S/H is only slightly above 1 for the minimum trace spacing achievable with the technology, which means that the technology capabilities are beyond what can be practically used by the product. By design, on some products, Pentium® 4 processors for example, the traces need to be spaced farther apart than the minimum spacing capability of the technology.

I/O Return Path

The final critical element of I/O interconnect design is the return path and the associated manifestations of a poor return path. "Return path" refers to the power and/or ground path that an I/O uses for reference, from which it draws current. This is a complex topic, and this paper cannot begin to do it justice. There are some best known methods (BKM's) and guidelines, which have been gathered as a result of experimentation and solutions of problems on recent products, which can be used to guide designers. Many of these are discussed in this section.

One manifestation of a poor return path is "SSO pushout." SSO pushout refers to the difference in timing between single-bit and multiple-bit switching. This is one of the items that system designers need to budget for up front in the design. Essentially, they need to identify the worst-case switching that the system will have to support and ensure that the design is robust enough to do this. Although there are many contributors to SSO pushout, recent validation work on the Pentium® II processor points to a large inductive return path as being a primary contributor.

Because it is difficult to model many return-path complexities, especially in the design phase when the layout is not frozen, a growing list of BKMs for avoiding potential problems is being formulated in the I/O interconnect design community at Intel. At the package level, the guidelines are as follows:

- Design a symmetric pinout with as low an I/O to power/ground pin ratio as possible within the pin budget for the package size. For example, in Figure 13, pinout A is better than pinout B because A is more symmetric.

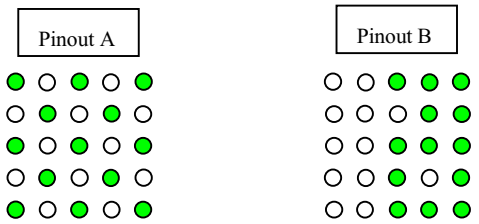

Figure 13: Pinout example

- Ensure that each signal has a clearly defined return path from the die to the land.
- Avoid voids or discontinuities in the return path.

Some general system BKM's, which should be used at each level of design, have also been developed [7]:

- Avoid signal lines that change reference planes, i.e., are referenced to ground in one part of the design and power in another part.
- Avoid referencing signal lines to power only, with no ground reference.
- Avoid transitioning from dual-referenced signals in one part of the design to single-referenced signals in another part of the design.
- Avoid referencing a signal to a power/ground reference that is not the primary reference for that signal. For example, there are often multiple power planes in a package. A signal must be referenced to the power plane that is supplying power for that signal not just to any power plane in the package.
- Avoid routing signal lines so that they cross splits in a reference plane.
- Avoid referencing lines to shielded lines that are floating or not adequately tied to power/ground with multiple vias.
- Avoid routing over voids in the plane, e.g., via voids in the ground plane.

POWER DELIVERY CHALLENGES AND SOLUTIONS

The challenges in package-level power delivery involve not only the package technology and design but also the selection and placement of discrete chip capacitors and of the voltage regulator module (VRM). Before discussing design practices, it is important to understand the modeling and simulation process that must be used in order to correctly design the package-level power delivery network. There are four main components that must be properly modeled: the package-level interconnect, discrete capacitors, the die power map, and the VRM [8].

The Package-Level Interconnect

The package-level interconnect is modeled using a 3D distributed, integrated, lumped model. It is important to obtain an accurate package model in order to enhance prediction of the whole power system performance. This network is connected to the system motherboard through several inductive and resistive socket pin elements, which terminate at one node on the motherboard. The motherboard, represented by its L, R, and C parasitics, in turn connects to a simple VRM model consisting of the VRM L, R, and C filters and input voltage. Figure 14 shows a representation of the full 3D-circuit network for a power delivery system.

Using 3D integrated models for the power delivery network allows analysis of subtle design changes that can result in much improved power delivery. Because the model is a highly distributed model, it also allows for the study of different capacitor placement scenarios. This allows optimization of the package design and chip capacitor layout early in the design phase instead of during validation testing, which is more costly. It also prevents over-design of the power delivery solution, which is also costly.

An example indicative of the type of analysis used in designing the Pentium® II and Itanium™ packages is useful in illustrating some of the design methods that enable more effective power delivery designs. In order to reduce the total loop inductance of the package-level interconnect, high coupling between power and ground structures is desired. To achieve that, the power/ground planes are divided into multiple strips with alternating power and ground bussing. The direction of power/ground strips in the same layer needs to be

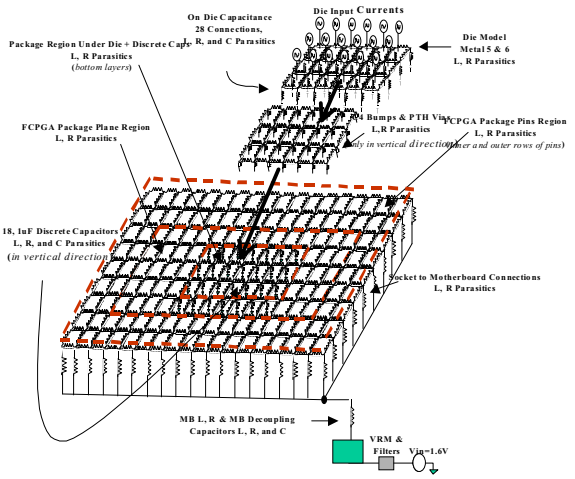

Figure 14: 3D network for a power delivery system

perpendicular to the power/ground strips in adjacent layers to enable via connection between layers. This is shown in the left of Figure 15.

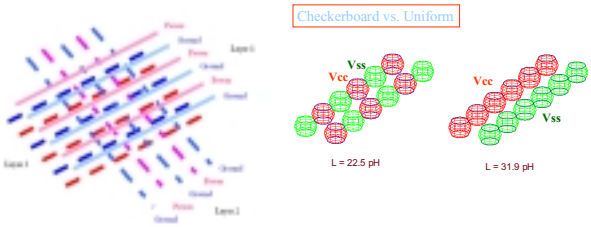

Figure 15: Multilayers of power strips and checkerboard vas

Discrete Capacitors

The next parts of the power delivery network are the discrete capacitors. Traditionally, a capacitor is modeled as a series resistance/inductance/capacitance (RLC) circuit. But the L (inductance) in the conventional RLC circuit is only attributed to the ESL of the capacitor itself. To accurately include the interaction between a capacitor, pads, vias, and power/ground planes, the inductance of the entire structure is computed using 3D electromagnetic modeling [9,10]. The capacitor itself is modeled as a metal sheet. Figure 16 shows an example of the 3D integrated model of a chip capacitor and its equivalent circuit model.

Due to the complexity of the problem, the integrated model of the chip capacitor is divided into two pieces to ease numerical model convergence. Table 1 shows the comparison of self inductance values of 0612, 0508, 0306, and 1206IDC capacitors and the effective inductance values of the integrated model for the capacitors with capacitor pads, vias, and package interactions.

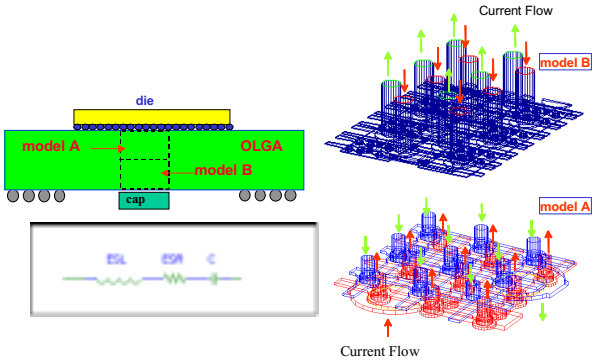

Figure 16: Capacitor mounted on a package and associated circuit model

Table 1: Self inductance vs. effective inductance

	0612	0508	0306	1206
Modeled Lself	310	300	155	57
Modeled Leff	220	205	166	72

The comparison shows that the mutual inductance between the capacitor itself, pads, power/ground planes, and power/ground buses can significantly reduce the total effective inductance of capacitors with 0612 and 0508 form-factors. For the 0306 capacitor, the inductance contributed from vias is 29 pH, and the ESL of the capacitor is 155 pH. The total sum of inductances of the vias and of the capacitor itself is 184 pH, which is larger than the effective inductance value of 166 pH obtained from the integrated model. The 1206 IDC results are similar.

The Die Power Map

It is important to obtain an accurate power map in order to predict the performance of a power delivery system. The on-chip interconnect model includes metal layers of the die that are combined to form a single 2D grid-like circuit network. The nodes of the 2D-circuit network are vertically connected to C4 bumps, which are in turn connected to the package circuit network underneath. In this 3D distributed circuit network, the C4 elements are represented by L and R components. The full die load model consists of MxN cells. For each cell, the die load model has an input current and on-die decoupling capacitance and resistance. Figure 17 illustrates a power map of a hypothetical die.

VRM

The final critical area in power delivery modeling and design is the VRM modeling and design [11]-[13]. Figure 18(a) represents the equivalent circuit of a VRM model. The topology is a buck type wherein the main

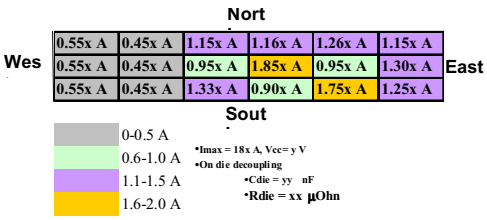

Figure 17: Die power map example

switch S_1 enables the step down of the input voltage. L_f and C_f constitute the low pass output filter elements. The duty ratio D of the converter is the ratio of the input voltage V_{in} and the output voltage V_0. Switch S_2 conducts at all other times when the main switch S_1 is off. Thus, the current I_L in the inductor L_f varies between values of I_{max} and I_{min} as shown in Figure 18(b). VRM switching frequency f_s pertains to the rate at which the two switches are switched. The switch S_2 can be a fast diode or a MOSFET. When the switch S_2 consists of a MOSFET instead of a diode, the configuration represents a synchronous rectifier topology. E_{ref} is the voltage whose value is to be maintained across the monitor points as V_{cc} in any power model. It serves as one of the inputs to a comparator inside the controller that implements both voltage and current feedback.

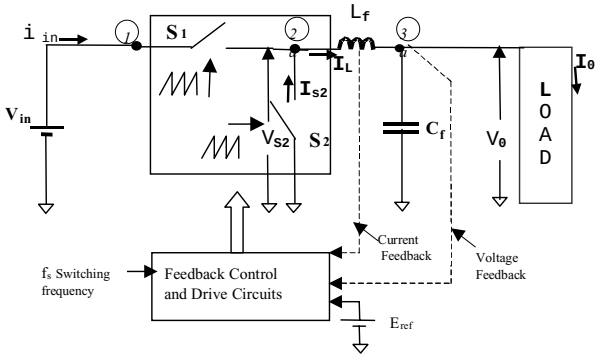

Figure 18(a): Schematic of a synchronous-rectifier based buck converter

Multiphase VRM's have become the norm in modern switching power supplies due to several advantages in terms of cost, size, and response characteristics achievable. Figure 19 shows the block diagram of a four-phase VRM without the detailed control and driver circuitry.

Extensive simulations have been done to study the VRM model performance. A few examples are given in the remainder of this section to illustrate some of the modeling and design decisions that must be made in designing the package-level interconnect power delivery system. The first example is a study with a simple VRM

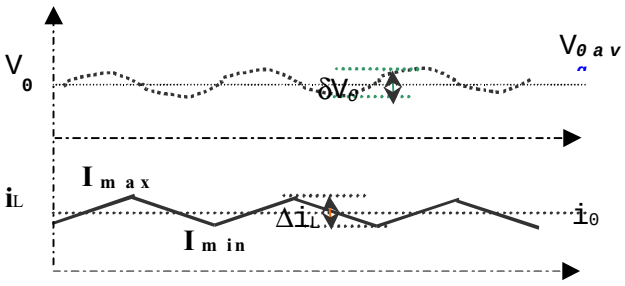

Figure 18(b): Voltage across VRM bulk capacitor Cf and the current in inductor Lf

model fed with a fixed voltage source. The simple VRM model consists of a simple battery in series with a small resistance and an inductance. A bulk capacitor is held across this combination. Then the droop study with the complex VRM model serving as the power supply to the model is presented. Finally, a comparison is made of the data when a single-phase and a four-phase VRM supply the input power to the circuit to investigate the effectiveness of each type of VRM.

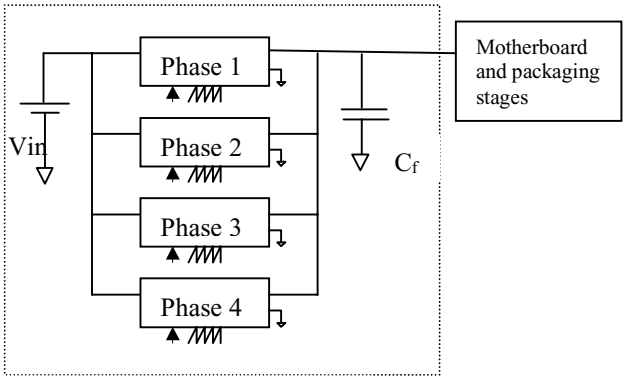

Figure 19: Four-phase VRM block diagram

VRM Modeling

For the study looking at the simple and complex VRM models, the waveform of the voltage V_{die} across the die for both the simple source and the four-phase VRM is illustrated in the bottom traces of Figure 20(a). The top trace is of the voltage V_{cc} measured after the power supply connector. The voltage droop target for both cases is 10% below V_{cc}. Three droops are expected in the waveforms corresponding to the package-die, interposer-package and VRM-interposer loops. In the waveforms

that pertain to the complex VRM model case, the VRM feedback control loop responds and the voltage begins to rise to the required V_{cc} value after the third droop. In the simple model case, the voltage continuously decreases as there is no control and hence a third droop is not seen. The first droop voltages are 1.133V and 1.130V, respectively, for the simple and complex models. The second droop voltages are 1.164V and 1.163V, respectively. The third droop voltage for the case with the complex VRM is 1.158V. This means that for the first and the second voltage droops, no significant difference between using a simple VRM model and a four-phase VRM is seen. A comparison of the transient response of the VRM that feeds the power delivery system is illustrated in Figure 20(b). The current supplied at the die and the current in the connector, which is the sum of all VRM phase currents, are shown. The current level achieved from a simple VRM is comparatively lower than the current level achieved from a four-phase VRM model, which explains the constantly decreasing voltage in the simple model in response to the load transient.

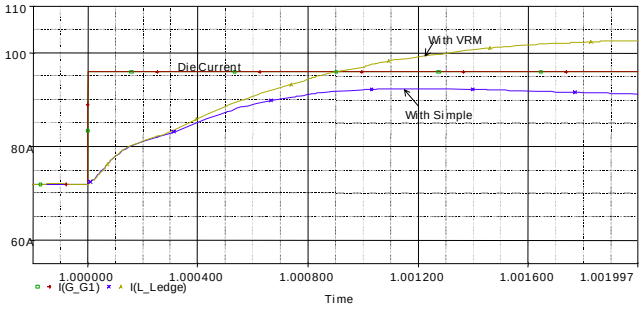

Figure 20(a): Waveforms of voltage and currents with simple source and VRM source – voltage across die and after POD connector

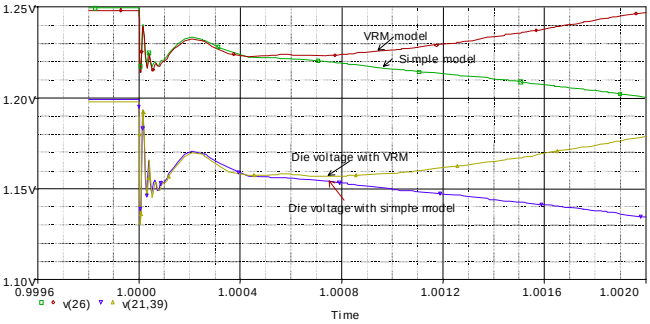

Figure 20(b): Waveforms of voltage and currents with simple source and VRM source – current transient at the die and the total VRM output current

VRM Design

Now that the need for using a complex VRM model has been established for accurately predicting the third-droop voltage, a study of the effectiveness of a single-phase vs. a four-phase VRM will be shown. A 250kHz single-phase and four-phase 1MHz/ph VRM were tied as inputs to the lumped power model. The waveforms of the voltages across the die and after the power supply connector are shown in Figure 21(a). The currents at the output of the VRM and at the die are shown in Figure 21(b). The response time of the four-phase model is much smaller than that of the single-phase model due to its much smaller inductance per phase and the much smaller bulk capacitance. Similar observations can be made regarding the current response. The single-phase VRM provides a slower and flatter transient response to the load compared to the four-phase VRM. Also a comparison of the droop voltages obtainable in the four-phase and single-phase cases is shown in Table 2, in which the ideal case entries are from the study on model complexity impact. From the second droop data in Table 2, we can observe that the single-phase VRM design is not as effective as the four-phase VRM design for this power delivery system.

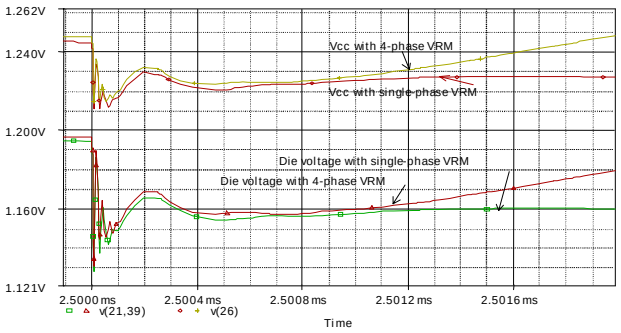

Figure 21(a): Waveforms of voltage and currents with single and four-phase VRM source – voltage across die and after POD connector

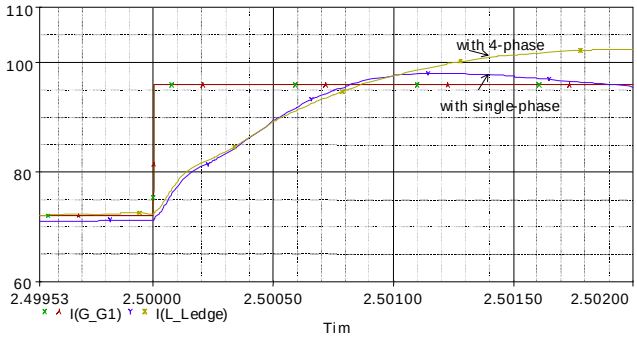

Figure 21(b): Waveforms of voltage and currents with single and four-phase VRM source – current transient at the die and the total VRM output current

Table 2. Effect of single and four-phase designs on droop characteristics

Case	Droop1(V)	Droop2 (V)	Droop3(V)
Ideal voltage source	1.133V	1.164V	NA
Single-phase VRM	1.130V	1.158V	1.154V
Four-phase VRM	1.130V	1.163V	1.158V

EMI CHALLENGES AND SOLUTIONS

Modeling the EMI characteristics at the package level is very difficult because EMI depends upon many different aspects of a system whose interactions are very complex. Because of the difficulty in building accurate EMI models for the package-level interconnect, much development work in the area of EMI solutions is experimental using test vehicles designed to mimic the product design. Some modeling is done using 3D full-wave, electromagnetic field solvers. Both the experimental and modeling work is usually of a qualitative rather than quantitative nature. The three areas that are usually investigated during development of an EMI solution are the EMI reduction properties of absorbers and shields added to the design; the most effective manner for grounding large metallic structures, such as thermal plates and heat sinks, so that they do not radiate and create EMI problems; and mitigation of radiation from high-speed traces in the package.

The design and development of the EMI solution for one of the Pentium® II processors can serve as an illustration of the types of EMI solution decisions that must be made at the package-level interconnect [14]. The Pentium® II processor server line of products was packaged on a cartridge with an enclosure. To aid in suppressing overall system-level EMI, the package designers decided to add an EMI shield to the cartridge enclosure. Several different materials were tested using test vehicles and benchmarked against a reference design with no shield. Testing was performed in a shielded EMI chamber, and the amount of reduction in EMI due to the package/cartridge only was measured. Figure 22 shows the results. Based upon the experimental results and consideration of cost and manufacturability, a suitable coating was chosen.

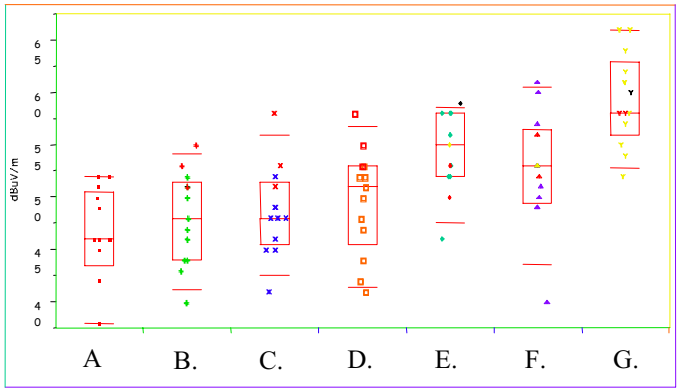

Figure 22: Emittance (x10) as a function of material tested for the frequency range of 140-1960 MHz.

A REVIEW OF RECENT DESIGNS TO MEET INTEL'S MARKET SEGMENTATION TARGETS

Tantamount to all package technology development and design activities is the need to meet cost targets while not compromising performance. Intel has increasingly segmented its microprocessor product lines in order to target different markets. This presents yet another challenge during the technology development and design cycle. A review of two recent package designs, one for Intel's Celeron™ product line, which is targeted to the low-cost market, and one for Intel's Itanium™ line, which is targeted to the high end of the market, is useful in illustrating the tradeoffs made to meet both cost and performance targets.

Celeron™ is Intel's value PC market line, meaning that maintaining a low cost is critical. Enabling good

performance is still an issue. When the core speeds for Celeron™ approached 500 MHz and front-side bus speeds increased to 100 MHz, it became necessary for Intel to migrate from its low-cost wirebond packaging to flip-chip packaging in order to maintain performance [15]. Flip-chip packaging was historically very expensive, however, which was contradictory to Intel's pricing strategy for the Celeron™ line. An effort to develop a low-cost flip-chip package interconnect ensued. The result was a flip-chip package with much coarser features than the high-end, high-cost flip-chip packaging that Intel uses for it higher end products. Following the design practices outlined in this paper, i.e., ensuring good signal return paths and keeping power loop inductances small, performance goals were met at a low cost.

Another example, taken from the design of the Itanium™ line of products, illustrates the tradeoffs in cost that can be made to achieve very high performance. Since Itanium™ is positioned as a high-end server product, cost requirements are not as stringent as for the Celeron™ line. Whereas, the Celeron™ package designs are limited to 4-6 layer designs, Itanium™ uses a 10-layer design and finer design rules to achieve higher performance [16]. Extra layers and finer design rules come at a cost hit; however, this is acceptable since performance is the most critical aspect in the design of this line of products. Figure 23 illustrates some of the analysis that went into determining that a 10-layer design should be used for Itanium™. The analysis results show a clear performance gain if more package layers and capacitors are used.

SUMMARY AND CONCLUSIONS

As microprocessor system performance continues to improve at exponential rates, the role of the package-level interconnect in enabling that performance will continue to increase in importance. Careful electrical analyses to ensure robust package designs that address signal integrity, power delivery, and EMI issues are critical to ensure that the package interconnect enables performance instead of inhibiting it. Intertwined in this design effort is the need to control the cost of the package interconnect solution. In addition, after careful review of the types of analyses that must be performed to ensure successful package designs, it is obvious that the package interconnect cannot be designed independent of the rest of the system. An integrated, holistic approach to system design in which performance and cost trade offs are made throughout the system in the appropriate and most convenient parts of the system is quickly becoming the most effective design approach. This approach will ensure that microprocessor performance trends are maintained.

ACKNOWLEDGMENTS

The analysis, design methods, and design examples discussed in this paper are the culmination of years of work by more people than the authors can properly acknowledge in this small section. We do explicitly thank PR Patel, the electrical group manager in ATD, who has continued to support and encourage our endeavors to improve package electrical design methods. We also acknowledge Dave White and Bob Sankman, the ATD package design managers, for their continuous support and recognition of the importance of good electrical design practices. We also thank our colleagues in the ATD electrical group for contributing in no small way to the body of knowledge and expertise outlined in this paper. Finally, we thank Adam Birr and Jung Kang for their thorough review of this paper and for their suggestions, which helped improve the paper.

REFERENCES

[1] Balanis, Constantine, *Advanced Engineering Electromagnetics*, New York, John Wiley & Sons, 1989.

[2] Paul, Clayton, *Introduction to Electromagnetic Compatibility*, New York, John Wiley & Sons, 1992.

[3] Johnson, Howard, and Martin Graham, *High-Speed Digital Design: A Handbook of Black Magic*, New Jersey, Prentice-Hall, Inc., 1993.

[4] Liao, Samuel, *Microwave Circuit Analysis and Amplifier Design*, New Jersey, Prentice-Hall, Inc., 1987.

[5] Wood, Dustin, "A Unique Degassing Hole Pattern to Minimize Impedance Variations in a OLGA Strip-line Package," *Intel Assembly & Test Technology Journal*, vol. 2, 1999.

[6] Vodrahalli, N., "OLGA2/INT2 Technology Target Spec.," Intel internal document, July 1999.

[7] Schoenborn, Zale, "Return Path Modeling: Preliminary Look at Identifying, Estimating, and Incorporating Effects into Willamette Sensitivity Analysis," Intel internal document, May 1998.

[8] Li, Y.L., David G. Figueroa, Shamala Chickamenahalli, Chee Yee Chung, Teong Guan Yew, Michael D. Cornelius, and Huong T. Do, "Enhancing Power Delivery System Through 3D Integrated Models, Optimized Designs, and Switching VRM Model," IEEE Transactions on Advanced Packaging to be published Aug. 2000.

[9] Li, Y.L., D.G. Figueroa, J.P. Rodriguez, L. Huang, J.C. Liao, M. Taniguchi, J. Canner, and T. Kondo, "A New Technique for High Frequency Characterization of

Capacitors," *Proceedings 48th Electronic Components and Technology Conference*, pp. 1384-1390, May 1998.

[10] Li, Y.L., D.G. Figueroa, T.G. Yew, C.Y. Chung, "Validation of integrated capacitor-via-planes model," *Proceedings of Electrical Performance of Electronic Packaging*, pp. 121-124, 1999.

[11] Zhang, Michael T., M. Jovanovic, F.C. Lee, "Design Considerations for Low-Voltage On-Board DC/DC Modules for Next Generations of Data Processing Circuits," IEEE Transactions on Power Electronics, vol. 11, no. 2, pp.328-337, March 1996.

[12] Zhou, Xunwei, Xingzhu Zhang, Jiangang Liu, Pit-Leong Wong, Jiabin Chen, Ho-Pu Wu, Luca Amoroso, FC. Lee and Dan Y. Chen, "Investigation of Candidate VRM Topologies for Future Microprocessors," *Applied Power Electronics Conference and Exposition*, APEC '98, pp. 145-150.

[13] Wong, Pit-Leong, FC. Lee, Xunwei Zhou, Jiabin Chen Xingzhu Zhang, "VRM Transient Study and Output Filter Design for Future Processors," *Applied Power Electronics Conference and Exposition*, APEC '98, pp. 410-415.

[14] Houle, Sabine, Kim Merley and Ian Roe, "EMI Solutions for Slot 2 Processors – Performance, Product Development – Risk and Cost," *Intel Assembly & Test Technology Journal*, vol. 2, 1999.

[15] Chung, Chee-Yee, Alex Waizman, Greg Taylor, Ananda Sarangi, Chris Baldwin, Bob Sankman, Altaf Hasan, "Flip Chip Pin Grid Array (FCPGA) Package, *Intel Design and Test Technology Conference*, to be presented August 2000.

[16] Zu, Larry, Tuanh Nguyen, Huong Do, and Lesley Polka, "McKinley Package Design: A Global Optimization Approach to Improve Electrical Performance," *Intel Assembly & Test Technology Journal*, vol. 2, 1999.

AUTHORS' BIOGRAPHIES

Lesley Polka joined Intel in 1994 and has worked since then in the Assembly Technology Development division, Chandler, AZ, on package electrical analysis, design and technology development. She has worked on several aspects of package development and design, including EMI analysis, development of measurement and modeling methods and support of product designs. Lesley presently leads a team focused on roadmap development for package-level I/O interconnects and on technology validation. Her technical interests include analysis and design of high-speed interconnects and package technology development. She obtained her B.S., M.S., and Ph.D. degrees, all in Electrical Egineering, from Arizona State University. Lesley can be reached at lesley.a.polka@intel.com.

Shamala Chickamenahalli received her B.E. and M.Tech. degrees in Electrical Engineering from Bangalore University and the Indian Institute of Technology, Khargapur, India, in 1983 and 1986, respectively. She obtained her Ph.D. degree in Electrical Engineering from the University of Kentucky in May 1995. From 1986 to 1988 Shamala worked as a senior engineer at Kirloskar Electric Co., India, and from 1988 to 1990 as a scientist at a defense research laboratory in India. She was an assistant professor from August 1994 to May 1999 at Wayne State University, Detroit, Michigan, where she taught and carried out DSP-based digital control for automotive applications. Currently, she is a senior packaging engineer at Intel Corporation, Chandler, AZ, where she concentrates on voltage regulator and power delivery schemes for microprocessors. Her research interests include modeling, analysis and simulation studies of novel low-power converter topologies. Shamala can be reached at shamala.chickamenahalli@intel.com.

Chee-Yee Chung obtained his M.S.E.E. degree from the University of Kentucky in 1995. He joined Intel Malaysia in 1996 and currently is a member of ATD Design, Chandler, AZ. While at Intel, Chee-Yee has worked on the PBGA, PLGA, OLGA and, most recently, the FCPGA packages. Chee-Yee's technical interests lie in all areas of high-speed interconnect electrical performance. Chee-Yee can be reached at chee.yee.chung@intel.com.

David G. Figueroa received his B.S. degree in Electrical Engineering from the University of Texas at El Paso in 1995. He joined Intel Corporation in Chandler, AZ, in 1996, where he has worked on the electrical modeling and characterization of interconnects and discrete components. His interests include power delivery, signal integrity, and the experimental characterization of interconnects and discrete components. David can be reached at david.g.figueroa@intel.com.

Yuan-Liang Li received his Ph.D. degree in Electrical and Computer Engineering from the University of Illinois at Urbana-Champagne in 1991. He worked for the U.S. Army Research Laboratory, Champaign, IL, from 1991 to 1995. In 1996, he joined Intel Corporation, ATD/DPD/Electrical Analysis and Design, Chandler, AZ, where he is currently leading a team working on research and development in the areas of signal integrity, power delivery, and EMI. His professional interests include electromagnetic modeling, simulation, and measurement of interconnects. He also serves as an adjunct associate

professor of Electrical and Computer Engineering at the University of Arizona, Tucson. Yuan can be reached at y.l.li@intel.com.

Kim Merley joined Intel in 1998 after 17 years experience in electronic design relating to EMC, ESD, and lighting protection at four different companies. He has worked on developing solutions to enable units to comply with EMI, ESD and lightning specifications. His work has included design of filters, EMI suppression components, shielding, bonding, gasketing, circuit board layouts, and power decoupling in addition to investigating many other aspects of EMC. He is currently involved with designing and testing test vehicles for EMI performance, and in research for better, less expensive, more reliable EMI suppression methods. He received his B.S.E.E. degree from Montana State University in 1980. Kim can be reached at kim.r.merley@intel.com.

Dustin Wood obtained his M.S.E.E degree from Arizona State University in 1997. He joined Intel in 1998 and has worked since then in the Assembly Technology Development Division in Chandler, AZ, on package electrical analysis, design, and technology development. He has worked on OLGA2, FCPGA1 and FCPGA2 package technologies and has focused primarily on signal integrity, advanced technologies, and packaging implementation at the product level. Dustin can be reached at dustin.p.wood@intel.com.

Larry Zu is a signal integrity engineer at Intel. Last year, his group supported Intel's leading-edge microprocessor package designs for iA32 and iA64 products. Over the past seven years at Intel and other companies, he has worked on several microprocessor package designs for desktop and server products for the Alpha™, StrongArm™, and Itanium™ microprocessors. He has a Ph.D. degree from Rutgers University in Electrical and Computer Engineering. Larry can be reached at larry.zu@intel.com.